

12 **EUROPEAN PATENT APPLICATION**

21 Application number : **94305231.6**

61 Int. Cl.<sup>6</sup> : **G06F 11/30**

22 Date of filing : **15.07.94**

30 Priority : **27.07.93 US 97506**

43 Date of publication of application :  
**01.02.95 Bulletin 95/05**

84 Designated Contracting States :  
**DE FR GB**


71 Applicant : **International Business Machines Corporation**  
**Old Orchard Road**  
**Armonk, N.Y. 10504 (US)**

72 Inventor : **Chang, David Yu**  
**7004 Windridge Cove**  
**Austin, Texas 78759 (US)**

74 Representative : **Litherland, David Peter**  
**I B M (UK) Ltd**  
**UK Intellectual Property Department**  
**Hursley Park**  
**Winchester, Hampshire SO21 2JN (GB)**

64 **Process monitoring in a multiprocessing server.**

67 A system and method for determining and displaying the status of client application programs executing on a multiprocessing server. Server process control blocks (5) and synchronization object descriptors (10) are created in the shared memory (11) of the server (3). Application program interfaces APIs (8) are linked to the control blocks and descriptors during the execution of the various multiprocessing application programs (7). A status utility (13) related to the service process monitor (12) selectively accesses information from the control blocks (5) and descriptors (10) to determine the status of the individual multiple processes executing on the server workstation (3). In a preferred form, the status information is conveyed to and displayed on a video display (4) associated with the service process monitor. In contrast to operating system monitors which disclose the status of all processes as a whole, the present server process monitor particularizes the information to the specific client process. Thereby, the information is of a granularity to identify processes which are hung up on semaphores, message queues, or the like. The information is at the level used by a system administrator or software developer.



The present invention relates in general to systems and methods for monitoring the activities of computers. More particularly, the invention is directed to a system and method for monitoring the states of client application type processes executing in a multiprocess server of a client-server network.

The client-server processing model has been widely adopted in the definition of distributed computing type networks. In the context of such networks, better performance with higher degrees of service concurrence have been exhibited by server operating systems which execute multiple client application programs through multiple processes. Examples are the OS/2® and AIX® operating system programs commercially available from IBM Corporation (OS/2 and AIX are registered trademarks of IBM Corporation). In contrast, single process operating systems require the server to await the completion of a current client's application program before commencing any aspect of a new client's application program.

The present concept of multiprocessing from the software perspective should not be confused with classical multiprocessing from the hardware perspective. From the hardware perspective, microprocessors such as the Intel Corp. models 80386 and 80486 incorporate time sharing features which accomplish multiprocessing through a time allocation for the different instructions being processed. In contrast, microprocessors such as the Intel Corp. model 80286 do not provide such a hardware capability, requiring that software manage any concurrent execution of multiple application programs. Operating system software which accomplishes this task for a 80286 type processor, and equally for the 80386 and 80486 microprocessors, is the aforementioned OS/2 operating system program. The present invention is directed to process management in the context of such an operating system, and not in the context of management by the microprocessor hardware.

The client-server network architecture is generally well known. With the advent of multiprocessing operating system capabilities in the servers, associating the activities occurring in the server to specific client application program processes has proven to be a significant challenge not only for the user clients but even for network administrators. Though operating systems, such as the aforementioned AIX program, provide resources for monitoring the state of a composite operating system on a server workstation, no contemporaneous information is provided about the states of the individual client application processes executing on the server. This level of information is particularly important to developers of client-server application programs. For example, presently available operating system monitors do not provide users with information regarding the server's work on a specific application program, or why a specific application program is hung up, or the identity of a sema-

phore delaying an application process. This deficiency is attributable to the fact that present operating system monitors do not link to the individual application processes, but rather, reflect the state of the composite of all server processes, viewed from the level of the operating system. Though trace log data could be generated in sufficient detail, the volume of the data requires storage to disk and time consuming analysis.

Therefore, a need exists for a monitoring system which provides contemporaneous information about the status of individual client application processes undergoing execution on a multiprocessing server in a client-server network.

Accordingly, the present invention provides a multiprocess server having a memory and operating in a client-server network and including a server process monitor comprising: means for creating multiple control blocks of the server process monitor in a part of the memory of the multiprocess server which is shared by the different server application processes; means for relating server processes to select control blocks; means for indicating the status of a process responsive to a server process monitor access of a control block.

Thus the present invention provides a system for monitoring individual server processes at the granularity of the client's application program, in a selective and contemporaneous fashion. Information regarding the status of each client application program as reflected by a server process is acquired and made available to the network administrator or client. In one form, the invention is directed to a monitoring system of a multiprocess server in a client-server network and comprises, a means for creating multiple control blocks of a server process monitor on a server processor, means for relating server application processes to control blocks in a server processor, means for storing the control blocks in memory shared by the different server application processes, and means for indicating the status of a server application process responsive to a server process monitor access of a control block. In another aspect, the invention relates to a method for accomplishing such application process selective monitoring.

In a preferred form, the invention involves a server process monitor program for creating control blocks and descriptors in a shared range of a server workstation memory. The control blocks are linked to and accessed by every server process executing a client application. In addition, every process running on the server has associated therewith synchronization object descriptors, defining semaphore or message queue states, which are similarly stored in shared memory. The control blocks and descriptors are registered with the server process monitor program upon creation and are accessed by the server process monitor program to derive, and subsequently

indicate on a video display or the like, the status of one or more of the server processes which are executing client applications.

The server process monitor function can also be performed by an operating system daemon process, which daemon process periodically reads the server process status information from the control blocks or descriptors in shared memory and displays the status on a dedicated window of a video display.

Irrespective of the particular form chosen, the information provides network administrators, software developers or field engineers with management, performance and diagnostic information at the granularity of each client process within a complex client-server network.

A preferred embodiment of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a schematic diagram depicting a client server network;

Figure 2 is a schematic diagram depicting server processes and the server operating environment; Figure 3 is a schematic diagram depicting a composite flow diagram for the processes of the system;

Figures 4-11 are schematics with the individual flow diagrams of those depicted in Figure 3;

Figure 12 is a schematic depicting the video operation display flow diagram.

The present invention is particularly useful in the context of a client-server network of the form depicted in Figure 1. Figure 1 shows a network 1 with a number of clients 2 and a server 3. Representative examples of the clients 1 and servers 3 would be PS/2 or RISC System/6000 workstations as are commercially available from IBM Corporation. A representative choice for network 1 in the context of PS/2 workstations would be Netbios and in such context include OS/2 Lan Server client code on respective workstations 2 and OS/2 Lan Server server code executing on server workstation 3. These operating systems are also commercially available from IBM Corporation. In the context of the RISC System/6000 workstation implementation, a preferred choice for the network would be TCP/IP and accordingly include in client 2 and server 3 workstations AIX type TCP/IP Server code, also commercially available from IBM Corporation.

Client workstations 2 transmit over network 1 requests that server 3 execute certain application program code responsive to commands issued by the client. In particular, the invention is directed to the server 3 executing in the multiprocessing mode of the aforementioned OS/2 or AIX operating systems, so that the various requests from the multiple clients timeshare the resources of server 3. The purpose of the server monitor is to determine and display, such as by way of video display 4, the status of each proc-

ess associated with each individual application program invoked by a respective client. This is in contrast to presently available operating system monitors which merely describe the overall state of the server and not the states of the individual processes. Those differences become crucial when a client, network administrator, software developer or field engineer needs to know the specific state of a client's process, not only in ascertaining its momentary status, but also in identifying, when, where and in what code and under what conditions process execution is temporarily or permanently interrupted.

Figure 2 schematically depicts the functional relationships between the processes and system elements needed to implement the present invention. Server workstation 3 executes multitasking operating system code 6, a code suitable to manage by software processes the relatively concurrent execution of application program code 7 and related application processes 8 of the multiple clients 2 (Figure 1) served by workstation 3. The present invention provides systems and methods for monitoring individual server processes at the granularity of the application program in contrast to the network or server operating system. Information regarding the status of individual server processes is acquired and made available for system administrator or user consideration.

An example of a multiprocessing server application program for which the present server process monitor has particular relevance is the IBM Parallel Database Manager Server program, commercially available from IBM Corporation. The Parallel Database Manager Server program is composed of multiple server processes, which include a pool of database manager agent processes, a deadlock detector process, a parallel database communication process, client communication processes, host gateway communication processes, and high availability processes. These server processes cooperate to provide the database manager service to the individual applications invoked by the clients. The process synchronization among the server processes is implemented through the process concurrence services of the base operating system, in this case the earlier noted OS/2 or AIX operating system. Semaphores, signals, and locks are examples of process concurrence services. The server processes also communicate among each other by sending information through shared memory or message queues.

These server processes can be forced into wait states, which states can be induced by a number of different reasons. For example, an agent process might wait for a reply message from a parallel database communication process, or for a new application request. Similarly, two agent processes might attempt to access the same database record at the same time, or a group of processes might be trapped into a deadlock situation. With so many server proc-

esses working concurrently in the system, a system level, but process specific, tool to monitor the server processes is needed by software developers, network administrators or service engineers when diagnosing malfunctions in a multiprocessing server.

The present server process monitor differs from the operating system process monitor. The process monitor of the base operating system describes the states of each process in terms of operating system parameters. For instance, the "PS" command in the AIX operating system causes the display of the user id, the process id, the parent process id, the start time, and the execution command of each AIX process. In contrast, the present server process monitor describes the server processes in terms of the state of each client's application process. Examples of valuable state information about the progress of an "agent" type process for the Parallel Database Manager Server are as follows:

- in free agent queue
- waiting on the database manager queue
- waiting on the parallel database queue
- processing database manager requests
- processing parallel database requests
- waiting on buffer queue services connection: token = xxxx, sid = x,
- waiting on closing buffer queue services connection: token = xxxx
- waiting on buffer distribution services message: token = xxxx, sid = x, rid = x
- waiting on buffer queue services data: token = xxxx, rid = x, rid = x
- waiting on fast communication manager memory request
- waiting on parallel database agent shared information
- waiting on table access: table token = xxxx
- waiting on access to data management services database control block
- waiting on access to data protection services database control block
- waiting to access to data protection services read buffer
- waiting to write a log
- deadlock detector waiting for time out
- waiting for log I/O done

From the examples of the states identified above it becomes apparent that the server process monitor provides state particulars about each individual server application process in contrast to merely identifying the presence of an application process.

Multiprocessing server 3 depicted in Figure 2 includes within its memory 9 a shared memory region 11. Server process control blocks 5 and synchronization object descriptors/blocks 10 are defined within shared memory 11. The placement of the control blocks and descriptors within the shared range of the memory addresses ensures that all the processes are

accessible to all of the control blocks of the process monitor. The common access also applies to server process monitor code 12, which defines a distinct server process status utility process 13. Control block and descriptor information is extracted and visually depicted on video display 4 by the utility process.

A server process is described in shared memory 11 by a server process control block. Such a block is created when a server process is generated and registered with the server process monitor. The server process control block preferably contains four fields:

proc\_id: the process id of the server process.

proc-type: by the nature of the server process, the server processes can be grouped into different process types. A server process can be a database agent process, communication process, a deadlock detector, et cetera. New process types can be created by the applications.

proc\_state: a process is either in "runnable" or "waiting" state.

syn\_obj\_handle: the handle of the synchronization object which associates with the server process. The handle of the synchronization object is the address of the synchronization object descriptor.

When a new server process type is created, a server process type record is also created. Each server process type record contains the following fields: proc\_type: the server process type identification.

proc\_desc: a text string that describes the function of the server process.

Synchronization objects such as latches, semaphores, wait post areas, or message queues are described by synchronization object descriptors in the server process monitor. Each synchronization object descriptor preferably contains the following data fields:

syn\_obj\_type: the types of the synchronization objects, including latch, wait post area, or message queue.

syn\_obj\_id: each synchronization object type has its own unique identifier. Latches are identified by latch handles, wait post areas are identified by wait post area handles, and message queues are referenced to message queue descriptors.

syn\_obj\_desc: a text string that describes the purpose of the synchronization object.

A set of application program interfaces suitable to use the data structures and described above is defined through a combination of a description, pseudo-code, and correspondence to a flow diagram of those depicted in the drawings.

The first application program interface (API) is to create a new server process type.

The input is:

proc\_desc: a text string that describes the function of the server process.

The output is:

proc\_type: the server process type identification.

Pseudocode defining the creation of a new server process type is as follows:

- create a new server process type.
- allocate a new server process type record.
- define the server process type identification.
- initialize the server process type record with the server process type identification and the server process description.
- return the server process type identification to the application program caller.

The flow diagram corresponding to the steps necessary to create a new server process type appears in Figure 4, which figure relates to the process composite in Figure 3.

After a new server process is created, it must register with the server process monitor. In that situation the input is:

- proc\_id: the process id of the server process.
- proc\_type: the server process type identification.

The registration process identified as reg\_svr\_proc has its output: ret\_sta: return status.

Pseudocode corresponding to the registration of a server process is as follows:

- create a server process control block for the server process.
- initialize the server process control block with the process id and the server process type.
- return control to the caller of the application program.

The flow diagram corresponding to the registration of a server process appears in Figure 5, which is likewise a part of the composite depicted in Figure 3.

The next application program interface (API) involves a change of the process type: chg\_svr\_proc. The change of the server process type from one to another involves an input of:

- proc\_id: the process id of the server process
- proc\_type: the new server process type identification.

The output of the application program interface is:

- ret\_sta: return status Pseudocode for changing a server process is as follows:
- update the server process control block of the server process with the new server process type.
- return control to the caller in the application program.

The corresponding flow diagram is depicted in Figure 6 of the drawings.

The application programming interface (API) reg\_syn\_obj registers a synchronization object such as a latch, semaphore, or message queue. The registration must be accomplished before it is referenced by a server process. The registration involves an in-

put of:

- syn\_obj\_type: the types of synchronization objects can be latches, wait post areas, or message queues. The synchronization object type identifications are defined by the server process monitor.
- syn\_obj\_id: each synchronization object type has its unique identifier. The synchronization object identifiers are defined by the base operating system when they are created.
- syn\_obj\_desc: a text string to describe the function of the synchronization object.

As an output the API provides:

- syn\_obj\_handle: the address of the synchronization object descriptor.

Pseudocode for implementing the API is set forth below in correspondence to Figure 7 of the drawings.

- Create a synchronization object descriptor.
- Update the synchronization object descriptor with a synchronization object type, synchronization object id, and the synchronization object descriptor.
- Return the synchronization object handle to the caller in the application program.

Before the server process calls the base operating system services to operate the synchronization object, the server process must call the wait\_syn\_obj to associate itself with the synchronization object.

The API involves an input of:

- proc\_id: the process id of the server process
- syn\_obj\_handle: the address of the synchronization object descriptor.

As an output of the API there is provided:

- ret\_sta: return status

Pseudocode corresponding to the flow diagram in Figure 8 of the drawings is set forth below:

- use the process id of the server process to locate the server process control block.
- change the process state of the server process control block from the: "runnable" to the "waiting" state.
- update the synchronization object handle of the server process control block.
- return control to the caller in the application program.

When the server process returns from the executing operations on the synchronization object, the server process calls run\_svr\_proc to change the server process state from "waiting" to "runnable".

The input is:

- proc\_id: the process id of the server process.

The output is:

- ret\_sta: return status.

The corresponding pseudocode, is depicted by flow diagram in Figure 9, involves the follows:

- Change the server process from "waiting" to "runnable" in its server process control block.
- Return control to the caller in the application

program.

A server process can be deregistered with the server process monitor by calling the `dereg_srv_proc` API. The server process control block of the server process will thereupon be freed. When a server process is terminated, by convention or otherwise, the server process exit routine calls `dereg_srv_proc` to deregister it from the server process monitor. The input to the API is:

--`proc_id`: the process id of the server process.

The output is:

--`ret_sta`: return status.

Pseudocode corresponding to the flow diagram in Figure 10 is as follows:

-- free the server process control block of the server process.

-- return control to the caller in the application program.

A synchronization object is deregistered from the server process monitor by calling an API identified as `dereg_syn_obj`. The `syn_obj_type` in the synchronization object in descriptor is changed to `invalid_obj`. The corresponding `syn_obj_id` in the synchronization object descriptor is changed to zero. The `syn_obj_desc` in the corresponding synchronization object descriptor is changed to a null string pointer.

The input to the API is:

--`syn_obj_handle`: the address of the synchronization object descriptor.

The output the of the API is:

--`ret_sta`: return status.

The flow diagram for this API appears in Figure 11 and corresponds to the following pseudocode:

-- change `syn_obj_type` in the synchronization object descriptor to `invalid_obj`.

-- change `syn_obj_id` in the synchronization object descriptor to zero.

-- change `syn_obj_desc` in the synchronization object descriptor to no.

-- return control to the caller in the application program.

An API utility suitable to convey server process status information to the video display, such as video display 4 in Figure 2, is presented by flow diagram in Figure 12. The server process status utility can be issued from any window of the base operating system. The utility spawns a process, process 13 in Figure 2, which has read access to the server process monitor residing in shared memory 11 of server workstation 3, as depicted in Figure 2. The utility reads the server process control block and synchronization object descriptor information and provides that information to video display terminal 4 in the format selected by the user. In a preferred form, the utility provides the user with options for selecting the server process status by process type, process state or the process id.

The utility includes resources for interpreting the synchronization object descriptor on which a process

is waiting in those situations where the server process is in a waiting state.

Pseudocode to display the server process status, corresponding to the flow diagram in Figure 12, is set forth as follows:

-- read all the server process control blocks.

-- if the server is runnable.

-- display the server process status.

-- else,

-- read the synchronization object descriptor.

-- endif,

-- return control to the caller in the server process status program.

Though the invention has been described and illustrated by way of a specific embodiment, the method and systems encompassed by the invention should be interpreted consistent with the breadth of the claims set forth hereinafter.

## Claims

1. A multiprocess server having a memory and operating in a client-server network and including a server process monitor comprising:
  - means for creating multiple control blocks of the server process monitor in a part of the memory of the multiprocess server which is shared by the different server application processes;
  - means for relating server processes to select control blocks;
  - means for indicating the status of a process responsive to a server process monitor access of a control block.
2. A server as claimed in claim 1, wherein the means for indicating the status of a server process comprises:
  - means for displaying the state of a server process using information from a control block accessed by the server process monitor.
3. A server as claimed in claim 1 or claim 2, further comprising:
  - means for registering a server process with the server process monitor upon creation of the server process in the multiprocess server.
4. A server as claimed in any preceding claim, further comprising:
  - means for relating synchronization objects of the server processes to descriptors accessible by the server process monitor.
5. A server as claimed in claim 4, further comprising:
  - means for registering a synchronization

object with the server process monitor upon creation of the synchronization object in the multiprocess server.

6. A method of operating a server process monitor of a multiprocess server having a memory and operating in a client-server network, comprising the steps of:  
    creating multiple control blocks of the server process monitor in a part of the memory of the multiprocess server which is shared by the different server application processes;  
    relating server processes to select control blocks; and  
    indicating the status of a server process responsive to a server process monitor access of a control block.
7. A method as claimed in claim 6, wherein the step of indicating the status of a server process comprises:  
    displaying the state of a server process using information from a control block accessed by the server process monitor.
8. A method as claimed in claim 6 or claim 7, further comprising the step of:  
    registering a server process with the server process monitor upon creation of the server process in the multiprocess server.
9. A method as claimed in any of claims 6 to 8, further comprising the step of:  
    relating synchronization objects of server processes to descriptors accessible by the server process monitor.
10. A method as claimed in claim 9, further comprising:  
    registering a synchronization object with the server process monitor upon creation of the synchronization object in the multiprocess server.

45

50

55

7





FIG. 1



FIG. 2



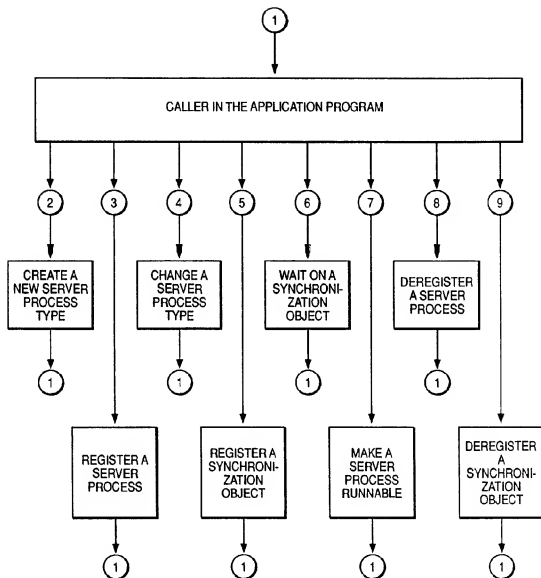


FIG. 3

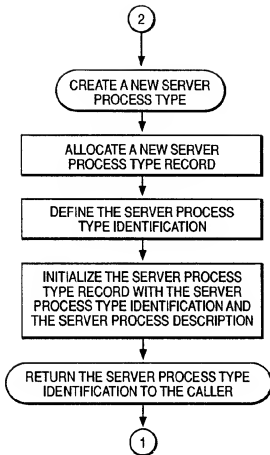


FIG. 4






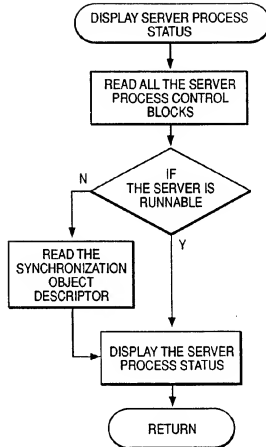




FIG. 5







European Patent  
Office

# EUROPEAN SEARCH REPORT

Application Number  
EP 94 30 5231

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (In CLC)
A	IBM TECHNICAL DISCLOSURE BULLETIN., vol.36, no.6B, June 1993, NEW YORK US pages 269 - 270 'Remote Database Status' * the whole document *	1-3,6-8	G06F11/30
A	IBM TECHNICAL DISCLOSURE BULLETIN., vol.31, no.8, January 1989, NEW YORK US pages 195 - 198 'Method for Managing Client/Server Relationships in the AIX Operating System' * the whole document *	1,6	
A	IBM TECHNICAL DISCLOSURE BULLETIN., vol.28, no.9, February 1986, NEW YORK US pages 3848 - 3855 'Multi-thread Host Server for PC Support' * page 3849, line 21 - page 3853, line 37 *	1,4,6,9	
A	AT&T 'UNIX System 5 (User's Reference Manual)', 1987, PRENTICE HALL, ENGLEWOOD CLIFFS, US PS(1) Command * page 210 - page 213 * -----		<div>TECHNICAL FIELDS SEARCHED (Int. Cl. 6)</div> <div>G06F</div>
The present search report has been drawn up for all claims			
Place of search <b>THE HAGUE</b>		Date of completion of the search <b>8 November 1994</b>	Examiner <b>Herreman, G</b>
<div>CATEGORY OF CITED DOCUMENTS</div> <div> <p>X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure F: intermediate document</p> <p>T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons</p> <p>&amp;: number of the same patent family, corresponding document</p> </div>			

EP 0 636 985 A1 (1994)